ABSTRACT OF THE INVENTION

A system for pipelined processing of program data uses a series of servers. A server is an event loop optimized for performing a particular task in the overall program. Servers perform their respective tasks in response to receiving work packets. A work packet is a data structure containing the data needed to perform that task. Each server has a queue of pending work packets. When there are work packets in a server's queue, the server performs the required function using the user context data contained in the packet, and may generate a reply code when complete. A server may run on a single CPU until all work packets in its queue have been processed, or it may also be shared among two or more CPUs. Each server performs its designated task for all users and thus, the general effect of this method is that each CPU performs a single task for all users, thereby insuring that the global context state for each task tends to remain in the cache of the CPU performing that task.